

CSCO - 7370 US P

UNITED STATES PATENT APPLICATION FOR

CONNECTION ROUTING BASED ON LINK UTILIZATION

Inventors:

Shawn Hsu
Yiwen Chen
Mana Palai

Prepared by:

WAGNER, MURABITO & HAO LLP
TWO NORTH MARKET STREET
THIRD FLOOR
SAN JOSE, CALIFORNIA 95113
(408) 938-9060

CONNECTION ROUTING BASED ON LINK UTILIZATION

FIELD OF THE INVENTION

Embodiments of the present invention relate to Asynchronous Transfer
5 Mode networking. More particularly, embodiments of the present invention
provide a method and system for connection grooming based on link utilization.

BACKGROUND

Asynchronous Transfer Mode (ATM) is a leading networking technology
10 for a wide variety of networking needs. ATM protocols are designed to handle
isochronous (time critical) data such as video and telephony, in addition to more
conventional data communications between computers or other networks. ATM
is a 'connection oriented' technology. Prior to any data transfer over an ATM
network, a connection (permanent or switched) is established which satisfies
15 the required data transfer quality of service (QOS), e.g., available bandwidth.
Such connections are frequently referred to as "calls."

ATM is particularly well suited to support virtual networking due to its
connection oriented nature. A capability to establish a private network between
20 any end systems provides physical location independence and privacy.
Location independence means that end users can be put on the same virtual
local area network (LAN) independently of their physical location. The Private
Network to Network Interface (PNNI) protocol, commercially available from the
ATM Forum of San Francisco, California, is generally used to establish virtual
25 connections across an ATM network.

A large ATM network is generally set up as a hierarchy of sets of smaller networks frequently referred to as "peer groups." For example, a low level peer group can comprise network nodes in a single metropolitan area. A higher
5 level network abstraction can comprise all such clusters as well as links between such clusters.

The term "routing" within ATM networks is generally understood to refer to or to describe a process of locating a target end system and selecting the
10 best possible path to reach the target. This function is similar to the routing provided by internet protocol (IP) networks. However, ATM routing has been extended to support two major requirements generally not available in IP networks, quality of service (QOS) routing and scalability. Quality of service routing supports classical data applications and real time applications such as
15 videoconferencing on the same physical infrastructure. Scalability supports very small networks of few switches up to world wide networks of thousands switches using a single protocol. Signaling is required due to the 'connection oriented' nature of ATM technology. When setting up (tearing down) a connection, each switch on the selected path needs to allocate (release) the
20 resources required to satisfy the QOS requirement of the connection. One such signaling message is the "setup" message, which is used to initiate a connection.

A virtual connection is often referred to as a Smart (or Soft) Permanent
25 Virtual Circuit (SPVC). SPVC is generally understood to refer to or to describe

a communications or communications medium which is “permanently” (for the life of the circuit) provisioned at the end points, but switched in the middle. An SPVC generally connects “point A” to “point B.” However, the actual routing path on the physical network is generally unknown to the user. Such virtual
5 circuits are usually long-lived, e.g., the life of such circuits is of the order of months.

Due to changing conditions on the underlying physical network, the best routing for a particular connection (virtual circuit) may change over time. For
10 example, assume that a direct path exists from point A to point B. However, if there is insufficient capacity available on that link for a new connection, the new connection can be routed from point A to point B via point C. At some later time, for example because of decreased usage of the direct path from point A to point B, sufficient capacity for the connection becomes available on the direct path
15 from point A to point B. A process known as connection “grooming” (also known as “route optimization”) periodically, e.g., once per day, can examine and optimize the path of a connection. If a routing change is to be made, generally a new route is established prior to tearing down the prior path. This is generally known as “soft” rerouting, and can utilize a “soft setup” signal or message.

20

During a connection grooming process, it is desirable to choose links that have the lowest “cost” for a connection. Many factors can be considered in determining a cost, including latency and delay for a link. Due to route optimization at call setup (initialization) and during grooming, low cost links tend
25 to fill up, e.g., reach capacity. Since the capacity of a virtual circuit is allocated

to that virtual circuit, a call tends to stay on the best, or lowest cost, link. As an unfortunate consequence, the “best” links are typically unavailable for new calls once the bandwidth for these links has been allocated to existing calls. In addition, any excess capacity is concentrated in the “worst” links. Therefore, an
5 interruption to the “best” links, e.g., a fault in such links, can cause virtual circuits to be rerouted from the “best” links to the “worst” links, which can have unacceptable quality of service.

Accordingly, it is desirable to route a virtual circuit according to the
10 utilization of a link. A further desire exists for communicating a link utilization factor through the links of a network. A still further desire exists for the above mentioned capabilities to be achieved with a scalable and cross-platform architecture that is compatible with industry standards.

SUMMARY OF THE INVENTION

Embodiments in accordance with the present invention route a virtual circuit according to the utilization of a link. Further embodiments in accordance with the present invention communicate a link utilization limit through the links of a network. Still further embodiments in accordance with the present invention achieve the above mentioned capabilities with a scalable and cross-platform architecture that is compatible with industry standards.

10 A method and system for connection routing based on link utilization are disclosed. A link utilization limit is accessed at a node of a network. The link utilization limit is compared to a utilization of a link coupled to the node. If the utilization of the link is greater than the link utilization limit, then the link is excluded from a virtual circuit.

15

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a block diagram of an exemplary network, upon which embodiments of the present invention can be implemented.

5

Figure 2 illustrates a flow chart for a method, in accordance with embodiments of the present invention.

Figure 3 illustrates a network switch apparatus, in accordance with
10 embodiments of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

In the following detailed description of the present invention, connection routing based on link utilization, numerous specific details are set forth in order to provide a thorough understanding of the present invention. However, it will be recognized by one skilled in the art that the present invention may be practiced without these specific details or with equivalents thereof. In other instances, well-known methods, procedures, components, and circuits have not been described in detail as not to unnecessarily obscure aspects of the present invention.

10

NOTATION AND NOMENCLATURE

Some portions of the detailed descriptions which follow (e.g., process 200) are presented in terms of procedures, steps, logic blocks, processing, and other symbolic representations of operations on data bits that can be performed on computer memory. These descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. A procedure, computer executed step, logic block, process, etc., is here, and generally, conceived to be a self-consistent sequence of steps or instructions leading to a desired result.

15

The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated in a computer system. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

20

25

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated
5 otherwise as apparent from the following discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as "comparing" or "processing" or "computing" or "translating" or "calculating" or "determining" or "excluding" or "displaying" or "recognizing" or "generating" or "assigning" or "initiating" or "collecting" or "transferring" or "switching" or
10 "accessing" or "retrieving" or "receiving" or "issuing" or "measuring" or "conveying" or "sending" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented
15 as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

CONNECTION ROUTING BASED ON LINK UTILIZATION

20 Figure 1 illustrates a block diagram of an exemplary network 100, upon which embodiments of the present invention can be implemented. It is to be appreciated that embodiments in accordance with the present invention are well suited to a wide variety of network topologies. The circles of Figure 1 represent nodes of network 100 and the solid lines represent links between the
25 nodes indicated. Peer group 120 comprises nodes 121, 122, 123 and 124.

Peer group 130 comprises nodes 131, 132 and 133. Link 110 represents the coupling between peer groups 120 and 130. In actuality, peer group 120 is coupled to peer group 130 by links 140 and 150. Links 140 and 150 can also be sometimes known as trunk links, or trunks.

5

Within peer group 120 there are links coupling the various nodes of the peer group. Link 125 couples node 121 with node 124. Link 126 couples node 121 with node 123. Link 127 couples node 123 with node 122. Link 128 couples node 121 with node 122. Similarly, there are links that couple the
10 various nodes of peer group 130. Links 136 and 137 couple node 132 with node 133. Link 138 couples node 131 with node 133. Link 135 couples node 131 with node 132.

In general, each link of network 100 can have a different information
15 carrying capacity and cost associated with carrying information. The cost of carrying information can comprise many factors, for example, the transmission delay associated with the link. Network routing generally attempts to determine a least costly routing to establish a virtual circuit.

20 Utilization of a link generally refers to an amount of information carried by a link relative to the link's maximum capacity. Link utilization can also refer to an amount of a link's information capacity that has been reserved, e.g., for virtual circuit connections, relative to the link's maximum capacity.

Assuming available bandwidth, the best route from node 121 to node 133 can be from node 121 over link 140 to node 132 and then over link 137 to node 133. However, in some cases it is desirable to keep the utilization of a link below a pre-determined maximum value. For example, keeping a link utilization below a pre-determined maximum value can allow greater network flexibility in responding to the creation of new virtual circuit connections and/or in response to equipment or linkage disruptions.

The utilization of a link can generally be known only for nodes within a common peer group. For example, node 132 can determine the utilization of link 138, for example by inquiring this information of node 131, but node 121 can not determine the utilization of link 138. Consequently, the initiating node of a virtual circuit connection can not generally control or determine whether a routing of the virtual circuit connection comprises a link exceeding a pre-determined maximum utilization.

In accordance with embodiments of the present invention, a "link utilization limit" can be specified as part of the setup message that begins routing of a virtual circuit connection. A link utilization limit generally specifies a maximum desired utilization, e.g., as a percentage of a link's maximum information carrying capacity, of a link that will be included as part of a virtual circuit connection. A link utilization limit can specify a maximum desired utilization for a link prior to, or including the required information capacity of the virtual circuit connection.

It is to be appreciated that the term "link utilization limit" as used herein does not, in general, refer to the maximum information carrying capacity of a link. The maximum information carrying capacity of a link is a property of a specific link. It is substantially static, generally changing only when

5 fundamental characteristics of a link change, for example, if the physical media of the link is changed. In contrast, a link utilization limit can be a property of a specific virtual circuit. In general it can be different for different virtual circuits, independent of the physical links that form such circuits. A link utilization limit as used herein is typically less than the maximum information carrying capacity

10 of a link. It is however to be appreciated that a link utilization limit can be numerically equal to maximum information carrying capacity of a link in accordance with embodiments of the present invention.

By indicating a requirement that a particular virtual circuit connection not

15 be routed over links exceeding the specified link utilization limit, nodes along an attempted path can determine if a particular link meets the specified link utilization limit, and consequently if that particular link is available for the virtual circuit connection.

20 For example, an initiator of a virtual circuit connection can determine that it is desirable to use only links at or below 75 percent of such links' maximum capacity in the virtual circuit connection. The 75 percent limit can include capacity for the virtual circuit connection. Alternatively, the 75 percent limit can be applied to the condition of a link prior to accounting for additional bandwidth

25 of the virtual circuit connection.

In accordance with embodiments of the present invention, such a link utilization limit can be included in setup message(s) in a variety of ways. One exemplary method is to embed a link utilization limit within the Generic

- 5 Application Transport Information Element in a setup message. The Generic Application Transport Information Element is defined by the ATM standards and can be forwarded undamaged through nodes that do not support a link utilization limit.

- 10 It is to be appreciated that it is possible that a virtual circuit connection can not be established due to such a link utilization limit. For example, there can be sufficient information capacity on links required to establish a virtual circuit connection. However, one or more of those links can be utilized in excess of the desired link utilization limit. In such a case, in accordance with
- 15 embodiments of the present invention, a virtual circuit connection originator can decide if it is more desirable to establish the circuit or to enforce the desired link utilization limit. If it is desired to establish the circuit, a link utilization limit can be omitted from the setup message(s). Alternatively, a link utilization limit can be set to 100 percent, allowing any link with remaining information capacity
- 20 sufficient for the circuit to be utilized.

- In accordance with another embodiment of the present invention, a link utilization limit can be enforced during connection grooming. Connection grooming (or route optimization) is performed periodically once a connection is
- 25 established. During connection grooming, a new route utilizing only links with

less capacity utilization than a specified link utilization limit can be attempted. If such a routing is successful, then the virtual circuit connection can utilize the new route. Generally, the “cost” of the new route in relation to the cost of the existing route is also considered in a decision to change routing.

5

In accordance with yet another embodiment of the present invention, link utilization can be included in a calculation of the “cost” of routing a virtual circuit connection. For example, in addition to calculating a routing cost comprising terms such as latency of a link, transmission time, number of packets queued for
10 transmission over a link and the like, link utilization can be included in such a routing cost calculation so as to generally increase a link cost for highly utilized links and/or lower a link cost for less utilized links. It is to be appreciated that there are many ways of including link utilization in a routing cost calculation that are well suited to embodiments in accordance with the present invention.

15

In accordance with embodiments of the present invention, a routing cost comprising link utilization information can be calculated as follows:

$$\text{Link Cost}(\text{utilization}) = \text{link cost} * w * \text{utilization} \quad (\text{relation } 1)$$

20

Where “Link Cost(utilization)” is a link cost comprising link utilization information, “link cost” is a link cost that does not include link utilization, “w” is a weighting factor, and “utilization” is a ratio of actual utilization of a link to that link’s information capacity. Relation 1 generally reduces the utilization-based
25 link cost for links with a low utilization, beneficially favoring such links in routing.

Embodiments in accordance with the present invention are well suited to routing using a combination of utilization-based routing cost and a link utilization limit. For example, a virtual circuit connection can be routed (or a routing attempted) based on a lowest cost routing, where link utilization is considered in the cost computation. The routing can still be subject to an absolute link utilization limit. In this novel manner, low utilization links can be favored in routing a virtual circuit that does not comprise links utilized above a desired link utilization limit.

10

Figure 2 illustrates a flow chart for a method 200, in accordance with embodiments of the present invention. In block 210, a desired link utilization limit is included in a virtual circuit connection setup message. The link utilization limit indicates a maximum utilization for links to be used for the virtual circuit connection.

15

In block 220, the link utilization limit is accessed by a node of a network. In block 230, the link utilization limit is compared to the utilization of a link coupled to the node. The comparison of the link utilization limit can be to a current utilization of the link. Alternatively, the link utilization limit can be to a current utilization of the link plus an additional bandwidth required for the virtual circuit connection.

20

Responsive to the comparison of block 230, if utilization of the link is greater than the link utilization limit, in block 240 the link is not used for the virtual circuit connection.

5 Referring once again to Figure 1, to setup an exemplary virtual circuit from node 121 to node 133, a setup message can include a link utilization limit of 75 percent (block 210 of Figure 2). In accordance with embodiments of the present invention, a link utilization limit can be used by nodes of a network to limit inclusion of links having a utilization of greater than the link utilization limit
10 from a virtual circuit. One possible routing of such a virtual circuit is via links 140 and 137.

In accordance with block 220 of Figure 2, node 121 can access the link utilization limit of the setup message. In block 230 of Figure 2, the link utilization
15 limit can be compared to the utilization of link 140. Assuming that link 140 has a utilization of 80 percent, it will not be used in the present routing of the virtual circuit.

The virtual circuit can be routed, for example, via link 125 to node 124,
20 and then via link 150 to node 131, and then via link 138 to node 133, presuming that none of links 125, 150 nor 138 have a link utilization of 75 percent or greater.

Figure 3 illustrates a network switch apparatus 300, in accordance with
25 embodiments of the present invention. Network switch apparatus 300

comprises a processor 310, computer usable media 315 and switching fabric 320. Computer usable media 315 can comprise a variety of types of computer readable memory, e.g., read only memory (ROM), flash electrically erasable read only memory ("flash"), or random access memory (RAM) and the like. A
5 plurality of links, e.g., links 330 and 340, can be coupled to network switch apparatus 300. Links 330 and/or 340 can be a variety of network links, e.g., T3, E3, single or multi-mode optical fiber, coaxial cable and the like.

The general operation of network switch apparatus 300 is to receive a
10 communication, e.g., a packet, on one of a plurality of links and to switch the communication to another of the links. The switching can be based upon characteristics of the communication, e.g., an address within a packet, or the switch can be configured to direct information from one link to another link in a generally static configuration. Processor 310 can play a role in switching the
15 information among the links, however processor 310 is generally used to control switching fabric 320 and switching fabric 320 performs substantially all of the information switching function.

In accordance with embodiments of the present invention, computer
20 usable media 315 can comprise computer-readable program code to be executed by processor 310. Processor 310 can interpret a setup message requesting the initialization of a virtual circuit connection. If the setup message comprises a link utilization limit, processor 310 can compare the link utilization limit to a utilization of a link coupled to the network switch apparatus 300. If the

utilization of the link is greater than the link utilization limit, then network switch apparatus 300 will not use the link for the virtual circuit connection.

For example, if a setup message is accessed from link 340 comprising a
5 link utilization limit of 75 percent, link 330 will not be used in the virtual circuit connection if link 330 has a link utilization of 80 percent.

Embodiments in accordance with the present invention route a virtual
circuit according to the utilization of a link. Further embodiments in accordance
10 with the present invention communicate a link utilization limit through the links of a network. Still further embodiments in accordance with the present invention achieve the above mentioned capabilities with a scalable and cross-platform architecture that is compatible with industry standards.

15 Embodiments in accordance with the present invention, connection routing based on link utilization, are thus described. While the present invention has been described in particular embodiments, it should be appreciated that the present invention should not be construed as limited by such embodiments, but rather construed according to the below claims.